

Low rank factorization of the Coulomb integrals for periodic coupled cluster theory

Felix Hummel,^{1, a)} Theodoros Tsatsoulis,¹ and Andreas Grüneis^{1, 2, b)}

¹⁾Max Planck Institute for Solid State Research, Heisenbergstraße 1, D-70569 Stuttgart, Germany

²⁾Department Chemie, Technische Universität München (TUM) Lichtenbergstraße 4, D-85747 Garching, Germany

We study the decomposition of the Coulomb integrals of periodic systems into a tensor contraction of six matrices of which only two are distinct. We find that the Coulomb integrals can be well approximated in this form already with small matrices compared to the number of real space grid points. The cost of computing the matrices scales as $\mathcal{O}(N^4)$ using a regularized form of the alternating least squares algorithm. The studied factorization of the Coulomb integrals can be exploited to reduce the scaling of the computational cost of expensive tensor contractions appearing in the amplitude equations of coupled cluster methods with respect to system size. We apply the developed methodologies to calculate the adsorption energy of a single water molecule on a hexagonal boron nitride monolayer in a plane wave basis set and periodic boundary conditions.

Keywords: coupled cluster; Coulomb integrals; tensor rank decomposition; canonical polyadic decomposition

I. INTRODUCTION

The high dimensionality of the many electron wave function is one of the most limiting factors in applying highly accurate electronic structure theories to the solution of the many electron Schrödinger equation for real materials on an *ab initio* level¹. Many of the most widely used wave function based theories have a good balance between accuracy, computational cost and the number of parameters used to approximate the exact wave function. The efficiency is strongly affected by the computational complexity required to evaluate resultant expectation values.

Tensor rank decompositions (TRD) and low rank tensor approximations are ubiquitous in the field of electronic structure theory calculations. These techniques are essential to reduce the computational cost and memory footprints to calculate and store the approximate many electron wave function. Already the simplest level of approximation, Hartree–Fock (HF) theory can be regarded as a low rank approximation, employing an antisymmetrized outer product of one electron orbitals to approximate the full wave function. This low rank tensor approximation is identical to a single Slater determinant. However, HF theory neglects electronic correlation effects. Electronic correlation effects can be captured by extending the wave function basis with additional determinants. For this purpose, excited HF determinants can be employed. They are constructed by replacing occupied orbitals with unoccupied orbitals, forming a complete and orthogonal basis. Computationally, the basis of (excited) Slater determinants is very convenient. It introduces a large degree of sparsity to the full many electron Hamiltonian and simplifies the solution of the many

electron problem. Most entries in the sparse many electron Hamiltonian can be calculated directly from electron repulsion integrals. The memory footprint for the storage of these integrals in a canonical basis is very large and grows rapidly with respect to the number of orbitals. Therefore it is often necessary to calculate these integrals in a computationally efficient on the fly manner. The most widely used schemes for the calculation of electron repulsion integrals include^{2–6}: (i) prior calculation of the integrals in the employed atomic orbital basis and its subsequent transformation into a molecular orbital basis, and (ii) employing the resolution of identity approach. Computationally the resolution of identity approach is more efficient because it requires the calculation and storage of intermediate quantities with at most three indices. In passing we note that the expression of the integrals in terms of these intermediate quantities allows for rearranging nested summations in ring coupled cluster theories such that the scaling of the computational cost with respect to the system size can be reduced⁷.

Coupled cluster theory can also be viewed as a low rank tensor approximation to the exact configuration interaction wave function coefficients in the Slater determinant basis. The exponential ansatz used in coupled cluster theories effectively approximates the coefficients of highly excited determinants by outer products of cluster amplitudes with a lower rank. However, increasingly accurate levels of coupled cluster theories lead to increasingly steep polynomial scalings of the computational cost and memory with respect to the studied system sizes. In this work we seek to reduce the computational cost of coupled cluster theories without introducing additional approximations on the level of the employed wave function. This can be achieved by employing low rank tensor approximation techniques for the decomposition of the two electron integrals and the corresponding intermediate quantities obtained from the resolution of identity approach. Using the low rank decomposition, the nested summations in the amplitude equations of distin-

^{a)}Electronic mail: f.hummel@fkf.mpg.de

^{b)}Electronic mail: a.grüneis@fkf.mpg.de

guishable cluster singles and doubles theory (DCSD)^{8–10}, as well as of linearized coupled cluster singles and doubles theory can be rearranged such that the scaling of the computational cost is reduced from $\mathcal{O}(N^6)$ to $\mathcal{O}(N^5)$ without any further approximations. Additionally, we highlight that the low rank factorization of the Coulomb integrals also allows for reducing the scaling of the computationally most expensive terms in coupled cluster singles and doubles (CCSD) theory using a plane wave basis set.

We note that the methods outlined in this work share many similarities with other approaches that aim at the reduction of the computational cost in correlated wave function based theories. In particular we want to point out that the tensor hypercontraction (THC) technique introduced by Hohenstein *et al.* in Refs. 11–13 also performs a low rank tensor decomposition of the Coulomb integrals. Furthermore, a similar approach to the tensor rank decomposition method introduced in this work was discussed by Shen *et al.* in Ref. 14. In the work of Benedikt *et al.* it was shown that tensor rank decomposition techniques can even be applied to the decomposition of the coupled cluster amplitudes directly¹⁵. However, in contrast to the methods mentioned above we introduce an efficient numerical procedure to achieve the low rank tensor decomposition of the Coulomb integrals in periodic systems without the necessity of defining an *a priori* real space grid as it is the case for THC methods.

A. Structure of this work

The factorization of the Coulomb integrals tensor is obtained in two steps. In Section II we first discuss how the Coulomb integrals can be decomposed into a contraction of two third order tensors: $V_{sr}^{pq} \approx \Gamma_s^{*pF} \Gamma_{rF}^q$, where we refer to Γ_{rF}^q as optimized Coulomb vertex. Subsequently, we perform a Tensor Rank Decomposition (TRD) of the optimized Coulomb vertex into a contraction of three matrices: $\Gamma_{rF}^q \approx \Lambda_F^R \Pi_R^{*q} \Pi_r^R$. Section III describes the employed algorithms to compute this factorization.

Section IV outlines how this factorization can be employed by quantum chemistry methods and in Section V we study the application of the discussed approximations to different systems. Subsection V A focuses on the convergence of the TRD for total energies of the LiH solid, while we compute coupled cluster adsorption energies of water on the surface of a single BN sheet in Subsection V B.

B. Notation

We imply a sum over all free indices occurring at least twice within a product but nowhere else. We will use the letters i, j, k, l to label occupied spin orbitals, a, b, c, d to label virtual spin orbitals and p, q, r, s to label general spin orbitals. The letters R, S, T, U are used to denote el-

ements of the rank decomposition. The conjugate transpose of a tensor such as A_r^q is denoted by A_q^{*r} where lower and upper indices are swapped. Sequence numbers in iterations are given in superscript within parentheses, as in $A^{(n)}$. The Frobenius norm of a tensor A is denoted by $\|A\|$. Examples are:

$$V_{ij}^{ab} V_{ab}^{ij} = \sum_{a,b \in \text{virt.}, i,j \in \text{occ.}} V_{ij}^{ab} V_{ab}^{ij} \quad (1)$$

$$T_{ijk} - A_{iR} B_{jR} C_{kR} = T_{ijk} - \sum_{R=1}^{N_R} A_{iR} B_{jR} C_{kR} \quad (2)$$

$$\|\Gamma_{rF}^q\|^2 = \sum_{q,r,F} \Gamma_q^{*rF} \Gamma_{rF}^q \quad (3)$$

II. OPTIMIZED AUXILIARY FIELD APPROXIMATION

In this section we discuss how to approximate the Coulomb integrals, a tensor of fourth order, by a contraction of two considerably smaller tensors of third order: $V_{sr}^{pq} \approx \Gamma_s^{*pF} \Gamma_{rF}^q$, without actually calculating the entire tensor V_{sr}^{pq} .

Given the spin orbitals $\psi_q(\mathbf{x})$ from a Hartree-Fock (HF) or density functional theory (DFT) calculation, the (nonantisymmetrized) Coulomb integrals are defined by

$$V_{sr}^{pq} = \iint d\mathbf{x} d\mathbf{x}' \psi^{*p}(\mathbf{x}) \psi^{*q}(\mathbf{x}') \frac{1}{|\mathbf{r} - \mathbf{r}'|} \psi_r(\mathbf{x}') \psi_s(\mathbf{x}), \quad (4)$$

with $\mathbf{x} = (\sigma, \mathbf{r})$ and $\int d\mathbf{x} = \sum_{\sigma} \int d\mathbf{r}$. Owing to the translational invariance of the Coulomb kernel, we can separate the Coulomb integrals as follows

$$V_{sr}^{pq} = \int \frac{d\mathbf{G}}{(2\pi)^3} \Gamma_s^{*p}(\mathbf{G}) \Gamma_r^q(\mathbf{G}), \quad (5)$$

where the *Coulomb vertex* $\Gamma_r^q(\mathbf{G})$ is given by

$$\Gamma_r^q(\mathbf{G}) := \sqrt{\frac{4\pi}{\mathbf{G}^2}} \int d\mathbf{x} e^{-i\mathbf{G} \cdot \mathbf{r}} \psi^{*q}(\mathbf{x}) \psi_r(\mathbf{x}). \quad (6)$$

We let its discretization be $\tilde{\Gamma}_{rG}^q = \sqrt{w_G} \Gamma_r^q(\mathbf{G}_G)$ with the momentum grid points \mathbf{G}_G and the numerical integration weights w_G such that

$$\int \frac{d\mathbf{G}}{(2\pi)^3} \Gamma_s^{*p}(\mathbf{G}) \Gamma_r^q(\mathbf{G}) \approx \sum_{G=1}^{N_G} \tilde{\Gamma}_s^{*pG} \tilde{\Gamma}_{rG}^q. \quad (7)$$

The Coulomb vertex can be computed from the spin orbitals in $\mathcal{O}(N_p^2 N_G \log N_r)$ time employing a Fast Fourier Transform (FFT), where N_p , N_G and N_r denote the number of spin orbitals, momentum grid points and real space grid points, respectively. We note that the orbital overlap charge density $\psi^{*q}(\mathbf{x}) \psi_r(\mathbf{x})$ is approximated in the projector augmented wave method using Eq. (2.87) of Ref. 16 as implemented in the Vienna *ab initio* simulation package (VASP)^{17–19}.

In general, any two body operator can be split into a product of two single body operators coupled by an auxiliary field²⁰. In the case of the Coulomb interaction the auxiliary field has only one variable due to translational invariance, which is the momentum \mathbf{G} mediated by the interaction. Although $\tilde{\Gamma}_{rG}^q$ is in practice already a third order tensor, the number of momentum grid points N_G of the HF or DFT calculation is usually too large to continue with the tensor rank decomposition of the Coulomb vertex directly. A large set of momenta will have a small but nonnegligible contribution to the correlation energy and we seek a more compact set of auxiliary field variables with fewer relevant elements.

Let

$$\tilde{\Gamma}_G^I = U_G^F \Sigma_F^J W^{*I}_J, \quad \text{G} \begin{array}{|c|} \hline \tilde{\Gamma} \\ \hline \end{array} \text{I} = \text{G} \begin{array}{|c|} \hline U \\ \hline \end{array} \begin{array}{|c|} \hline \Sigma \\ \hline \end{array} \begin{array}{|c|} \hline W^* \\ \hline \end{array} \text{I} \quad (8)$$

be a singular value decomposition of the Coulomb vertex $\tilde{\Gamma}_G^I$, written as a matrix with the compound index $I = (q, r)$, where the singular values in Σ are sorted in descending order. Eq. (8) is also shown in form of a wiring diagram of the involved tensor contractions on the right. Taking only the largest $N_F < N_G$ singular values of the unapproximated Coulomb vertex $\tilde{\Gamma}$ into account we can define the *optimized auxiliary field (OAF) Coulomb vertex*

$$\Gamma_F^I := \Sigma_F^J W^{*I}_J = U^{*G}_F \tilde{\Gamma}_G^I. \quad \text{F} \begin{array}{|c|} \hline \Gamma \\ \hline \end{array} \text{I} := \text{F} \begin{array}{|c|} \hline U^* \\ \hline \end{array} \begin{array}{|c|} \hline \tilde{\Gamma} \\ \hline \end{array} \text{I} \quad (9)$$

Note that we write Γ without a tilde for the approximated Coulomb vertex, in contrast to usual convention, simply because we will not use the unapproximated vertex in any subsequent step.

We are only interested in the left singular vectors U_G^F associated to the largest singular values so we contract Eq. (8) from the right with $\tilde{\Gamma}_I^{*G'}$

$$\tilde{\Gamma}_G^I \tilde{\Gamma}_I^{*G'} = U_G^F \Sigma_F^J U^{*G'}_F =: E_G^{G'}, \quad (10)$$

transforming a singular value problem of a large $N_G \times N_p^2$ matrix into an eigenvalue problem of a comparatively small $N_G \times N_G$ hermitian matrix. The eigenvalues of E are the squares of the singular values of $\tilde{\Gamma}_G^I$, and the left eigenvectors of E associated to the largest eigenvalues are also the left singular vectors of $\tilde{\Gamma}$ we need in order to transform the Coulomb vertex $\tilde{\Gamma}_G^I$ into the optimized auxiliary field Coulomb vertex Γ_F^I according to Eq. (9). Note that this approach becomes numerically problematic for very small singular values since one only has access to their squares. However, we find that all N_F largest singular values needed for an accurate approximation of the Coulomb vertex are sufficiently large.

Inserting the singular value decomposition of the Coulomb vertex with sorted singular values from Eq. (8) into the discretized definition of the Coulomb integrals given in Eq. (7) yields a singular value decomposition of

the Coulomb integrals, also with sorted singular values:

$$V_J^I = W_J^K \Sigma_K^{2K} W^{*I}_K, \quad \text{J} \begin{array}{|c|} \hline V \\ \hline \end{array} \text{I} = \text{J} \begin{array}{|c|} \hline W \\ \hline \end{array} \begin{array}{|c|} \hline \Sigma^2 \\ \hline \end{array} \begin{array}{|c|} \hline W^* \\ \hline \end{array} \text{I} \quad (11)$$

where the Coulomb integrals V_{sr}^{pq} are now written in matrix form V_J^I with $I = (q, r)$ and $J = (s, p)$. Thus, using the optimized auxiliary field Coulomb vertex Γ_F^I instead of the full Coulomb vertex $\tilde{\Gamma}_G^I$ best approximates the Coulomb integrals with respect to the Frobenius norm of the difference:

$$V_{sr}^{pq} \approx \Gamma_s^{*pF} \Gamma_{rF}^q. \quad \text{p} \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \text{q} \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \approx \text{p} \begin{array}{|c|} \hline \Gamma^* \\ \hline \end{array} \text{s} \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \text{F} \begin{array}{|c|} \hline \Gamma \\ \hline \end{array} \text{r} \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \text{q} \quad (12)$$

III. DECOMPOSITION OF THE COULOMB VERTEX

The form of the Coulomb vertex in real space on the right hand side in Eq. (6) suggests that the optimized Coulomb vertex Γ_{rF}^q can be decomposed in an analogous manner into a product of three tensors of second order, denoted and depicted as follows

$$\Gamma_{rF}^q \approx \Lambda_F^R \Pi^{*q}_R \Pi_r^R. \quad \text{F} \begin{array}{|c|} \hline \Gamma \\ \hline \end{array} \text{r} \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \text{q} \approx \text{F} \begin{array}{|c|} \hline \Lambda \\ \hline \end{array} \begin{array}{|c|} \hline \Pi^* \\ \hline \end{array} \begin{array}{|c|} \hline \Pi \\ \hline \end{array} \text{R} \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \text{r} \quad (13)$$

We let N_R denote the number of *vertex indices* R and refer to it as the *rank* of the Coulomb vertex for a given quality of the approximation. We call the matrices Π_r^R and Λ_F^R *factor orbitals* and *Coulomb factors* of the Coulomb vertex, respectively.

The decomposition is invariant under scaling of the Coulomb factors Λ_F^R with any real scalar $a_R > 0$ while scaling the factor orbitals Π_r^R with a complex scalar c_R with $|c_R| = 1/\sqrt{a_R}$ for each value of R . One can also choose an alternative ansatz to Eq. (13) for approximating the Coulomb vertex Γ_{rF}^q which does not involve the conjugation of the factor orbitals on the outgoing index q . This ansatz reads

$$\Gamma_{rF}^q \approx \Lambda_F^R \Pi_q^R \Pi_r^R. \quad \text{F} \begin{array}{|c|} \hline \Gamma \\ \hline \end{array} \text{r} \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \text{q} \approx \text{F} \begin{array}{|c|} \hline \Lambda \\ \hline \end{array} \begin{array}{|c|} \hline \Pi \\ \hline \end{array} \begin{array}{|c|} \hline \Pi \\ \hline \end{array} \text{R} \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \text{r} \quad (14)$$

The above decomposition is invariant under scaling of the Coulomb factors Λ_F^R with any complex scalar $c_R \neq 0$ while scaling the factor orbitals Π_r^R with $\pm 1/\sqrt{c_R}$ for each value of R . This is in contrast to the symmetries of the ansatz according to Eq. (13). In the Alternating Least Square (ALS) approximation scheme, which we employ for fitting the factor orbitals Π and the Coulomb factors Λ , it is preferable to have the symmetries of Eq. (14)

since fixing one of the two factors removes all continuous symmetries from the other. This accelerates convergence and allows a smaller rank N_R in practice. A downside of this ansatz is that one loses the simple notion of particle and hole propagators as given in Eq. (34) and Eq. (35), respectively. The propagators can, for instance, be employed to calculate second order Møller–Plesset theory (MP2) correlation energies in $\mathcal{O}(N^4)$, as outlined in Subsection IV A. Unless otherwise stated we use the ansatz of Eq. (14) for the applications presented in Section V but we will continue to discuss the more widely applicable ansatz of Eq. (13).

If one chooses $N_R = N_r$, where N_r is the number of real space grid points, the validity of the ansatz follows directly from Eq. (6). We want to investigate how low N_R can be chosen compared to N_r for a sufficiently faithful decomposition having an error below 1% in the energies calculated from the factor matrices. For reference we use the error in the MP2 energy assuming that other terms, occurring for instance in the coupled cluster amplitude equations, exhibit a similar behavior. This accuracy is assumed sufficient since in practice only those terms will be calculated from the factor matrices that pose either computational or memory bottlenecks. Furthermore, we want to show that the ratio N_R/N_r for a sufficiently faithful decomposition is independent of the system size if the system is not too small.

A. Canonical polyadic decomposition algorithms

A factorization of a tensor according to the ansatz of Eq. (13) or (14) is referred to as canonical polyadic decomposition²¹ (CPD). For a given rank N_R , the factor orbitals Π_r^R and the Coulomb factors Λ_F^R can be fit by minimizing the square of the Frobenius norm of the difference

$$(\Lambda, \Pi) = \operatorname{argmin}_{\Lambda, \Pi} \left\| \Lambda_F^R \Pi^{*q} \Pi_r^R - \Gamma_{rF}^q \right\|^2. \quad (15)$$

The above optimization problem is high dimensional and nonquadratic. Conjugate gradient algorithms or other local algorithms may require thousands of steps until sufficiently converged. Global optimization algorithms try to tackle the problem by keeping a subset of the variables fixed and optimizing only the remaining variables. In the case of the alternating least squares²² (ALS) algorithm the optimization is done in turn over each matrix, while in the case of the cyclic coordinate descent²³ (CCD) algorithm the optimization is done in turn over each value of the index R . We have studied the performance of a regularized version of the ALS here.

B. Alternating least squares

In the case of three distinct factors $T_{ijk} \approx A_{iR} B_{jR} C_{kR}$ two of them can be regarded fixed leaving a least squares

problem for finding the optimal third factor. Each matrix is optimized in alternating order leading to the *alternating least squares* (ALS) algorithm

$$A^{(n+1)} := \operatorname{argmin}_A \left\| A_{iR} B_{jR}^{(n)} C_{kR}^{(n)} - T_{ijk} \right\|^2, \quad (16)$$

$$B^{(n+1)} := \operatorname{argmin}_B \left\| A_{iR}^{(n+1)} B_{jR} C_{kR}^{(n)} - T_{ijk} \right\|^2, \quad (17)$$

$$C^{(n+1)} := \operatorname{argmin}_C \left\| A_{iR}^{(n+1)} B_{jR}^{(n+1)} C_{kR} - T_{ijk} \right\|^2, \quad (18)$$

which has to be solved iteratively until convergence, starting with random matrices $A^{(0)}$, $B^{(0)}$ and $C^{(0)}$.

Each least squares problem has a unique solution, which can be written explicitly. For Eq. (16) it is for instance given by

$$A_{iR}^{(n+1)} = T_{ijk} B^{*jS} C^{*kS} G_{SR}^+, \quad (19)$$

omitting the iteration specification on B and C for brevity. G^+ denotes the Moore–Penrose pseudoinverse^{24,25} of the Gramian matrix G . For Eq. (16) the Gramian matrix is given by

$$G_{RS} = B^{*jS} C^{*kS} B_{jR} C_{kR}. \quad (20)$$

The expressions for the other matrices can be written in an analogous manner. When applying the ALS algorithm to decompose the Coulomb vertex the computationally most demanding steps are the calculation of the pseudo inverse G^+ scaling as $\mathcal{O}(N_R^3)$, as well as the contraction of T_{ijk} with either factor B^{*jS} or C^{*kS} in Eq. (19), depending on which is larger, scaling as $\mathcal{O}(N_p^2 N_F N_R)$.

C. Regularized alternating least squares

Although the ALS algorithm guarantees an improvement of the fit quality in each iteration the convergence can be very slow, especially when there are multiple minima for a factor A in different regions having all similar minimal values. In that case the best choice for A may vary strongly from iteration to iteration since updating the other factors B and C can change the order of the minima. This behavior is referred to as swamping²⁶ and it takes many iterations before the ALS algorithm converges to one region for each factor that globally minimizes the fit quality. Introducing a penalty on the distance to the previous iteration limits swamping and leads to the regularized ALS²⁷ (RALS) algorithm:

$$A^{(n+1)} := \operatorname{argmin}_A \left(\left\| A_{iR} B_{jR}^{(n)} C_{kR}^{(n)} - T_{ijk} \right\|^2 + \lambda_A^{(n)} \left\| A_{iR} - A_{iR}^{(n)} \right\|^2 \right), \quad (21)$$

$$B^{(n+1)} := \operatorname{argmin}_B \left(\left\| A_{iR}^{(n+1)} B_{jR} C_{kR}^{(n)} - T_{ijk} \right\|^2 + \lambda_B^{(n)} \left\| B_{jR} - B_{jR}^{(n)} \right\|^2 \right), \quad (22)$$

$$C^{(n+1)} := \operatorname{argmin}_C \left(\left\| A_{iR}^{(n+1)} B_{jR}^{(n+1)} C_{kR} - T_{ijk} \right\|^2 + \lambda_C^{(n)} \left\| C_{kR} - C_{kR}^{(n)} \right\|^2 \right). \quad (23)$$

The solution of each regularized least squares problem can again be given explicitly, here for instance for Eq. (21), and again omitting the iteration specification on B and C

$$A_{iR}^{(n+1)} = \left(T_{ijk} B^{*jS} C^{*kS} + \lambda_A^{(n)} A_{iS}^{(n)} \right) G_{SR}^+. \quad (24)$$

In the regularized case the Gramian G depends on the regularization parameter $\lambda_A^{(n)}$

$$G_{RS} = B^{*jS} C^{*kS} B_{jR} C_{kR} + \lambda_A^{(n)} \delta_{RS}, \quad (25)$$

where δ_{RS} denotes the Kronecker delta.

The regularization parameter $\lambda_A^{(n)}$ for finding $A^{(n+1)}$ in the n th iteration still remains to be determined. Too low values allow swamping to occur while too large values unnecessarily slow down the convergence. To estimate an efficient regularization parameter we assume that the fit quality $\|AB^{(n)}C^{(n)} - T\|^2$ in the term to minimize in Eq. (21) varies little from one iteration to the next. We also assume this for the local change of the fit quality with respect to each value in A . This allows us to relate the minimized term in the previous step $n-1$ to the minimized term in the step n for which we want to determine the regularization parameter:

$$\lambda^{(n-1)} \|A^{(n)} - A^{(n-1)}\|^2 \approx \lambda^{(n)} \|A^{(n+1)} - A^{(n)}\|^2. \quad (26)$$

If $A^{(n)}$ and $A^{(n+1)}$ have similar norm we can also relate their relative step sizes $s_A^{(n)}$ and $s_A^{(n+1)}$, by $\lambda^{(n-1)} s_A^{2(n)} \approx \lambda^{(n)} s_A^{2(n+1)}$, where the relative step size in the n th iteration is given by

$$s_A^{(n)} := \|A^{(n)} - A^{(n-1)}\| / \|A^{(n)}\|. \quad (27)$$

We want the relative iteration step size $s_A^{(n+1)}$ of the next iteration to be approximately as large as a chosen maximum value s_0 , which we refer to as *swamping threshold*. From that we define the estimated regularization parameter for the n th iteration

$$\hat{\lambda}_A^{(n)} := \lambda_A^{(n-1)} s_A^{2(n)} / s_0^2. \quad (28)$$

Using the above estimate directly results in a regularization which we find alternately too strong and too weak.

To ameliorate this we introduce a mixing of the estimated regularization parameter $\hat{\lambda}_A^{(n)}$ for the n th iteration, as above, with the regularization parameter $\lambda_A^{(n-1)}$ of the previous iteration to obtain the regularization parameter $\lambda_A^{(n)}$ employed for the n th iteration in the RALS:

$$\lambda_A^{(0)} := 1 \quad (29)$$

$$\lambda_A^{(n)} := \alpha \hat{\lambda}_A^{(n)} + (1 - \alpha) \lambda_A^{(n-1)}. \quad (30)$$

Regarding the choice of the swamping threshold s_0 and the mixing factor α we find that $s_0 = 1.0$ and $\alpha = 0.8$ offers a good compromise allowing quick convergence while still preventing swamping for the systems studied so far.

D. Quadratically occurring factors

In the case of the Coulomb vertex the factor orbitals Π_r^R occur quadratically. For finding the next estimate $\Pi^{(n+1)}$ in the alternating least squares algorithm we use an iterative algorithm similar to the Babylonian square root algorithm. Each subiteration is given by

$$\begin{aligned} \Pi^{(n+1,m+1)} &:= (1 - \beta) \Pi^{(n+1,m)} \\ &+ \beta \operatorname{argmin}_{\Pi} \left(\left\| \Lambda_F^{R(n+1)} \Pi^{*q(n+1,m)} \Pi_r^R - \Gamma_{rF}^q \right\|^2 \right. \\ &\quad \left. + \lambda_{\Pi}^{(n+1,m)} \left\| \Pi_r^R - \Pi_r^{(n+1,m)} \right\|^2 \right) \end{aligned} \quad (31)$$

with $\Pi^{(n+1,0)} := \Pi^{(n)}$ and the mixing factor $0 < \beta < 1$. Note that Π^* is a fixed parameter rather than a fitted one and that the regularization parameter $\lambda_{\Pi}^{(n+1,m)}$ needs to be determined for the m th subiteration similar to Eq. (30), however with $\lambda_{\Pi}^{(n+1,0)} = \lambda_{\Pi}^{(n)}$. The above iteration converges towards a solution of the quadratic problem. We use $\Pi^{(n+1,M)}$ and $\lambda_{\Pi}^{(n+1,M)}$ for the next estimate of $\Pi^{(n+1)}$ and $\lambda_{\Pi}^{(n+1)}$ in the RALS algorithm, respectively. The number of subiterations M needs to be sufficiently large, such that $\Pi^{(n+1)}$ is at least an improved solution of the entire fit problem compared to $\Pi^{(n)}$. A large number M of subiterations gives an estimate $\Pi^{(n+1)}$ that is close to the optimal choice of Π for a given $\Lambda^{(n+1)}$. However, the cost of each subiteration are similar to the cost of the fit of Λ in the RALS algorithm and as a good choice for minimizing the overall computational cost we find $\beta = 0.8$ and $M \geq 2$, but only as large such that the solution is an improvement. We point out that there are alternative methods for solving the quadratically occurring factors²⁸.

IV. APPLICATION OF THE LOW RANK FACTORIZATION

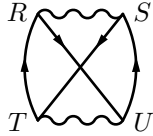
The algorithms described so far yield an approximate factorization of the Coulomb integrals of the form

$$V_{sr}^{pq} \approx \Pi_R^{*p} \Pi_S^{*q} \Lambda_R^{*F} \Lambda_F^S \Pi_r^S \Pi_s^R, \quad (32)$$

where the factors Π and Λ are $N_p \times N_R$ and $N_F \times N_R$ matrices, respectively. We find that the rank of the decomposition N_R is about an order of magnitude lower than the number of real space grid points of the original factors of the Coulomb integrals, being the orbitals $\psi_q(\mathbf{x})$. In Section V we study the convergence of the approximation in detail. In this section we discuss how this factorization can be applied to lower the scaling of the computational cost of wave function based methods such as second order Møller–Plesset (MP2) theory or coupled cluster theory.

A. MP2 from imaginary time propagators

The factorization of the Coulomb integrals permits evaluating the terms in the perturbation expansion by summing over all vertex indices R, S, T, \dots occurring in the term's diagram, contracting propagator matrices for each particle, hole and Coulomb line. For instance, the exchange term of second order Møller–Plesset (MP2) theory can be evaluated as follows:



$$= -\frac{1}{2} \int_0^\infty d\tau \sum_{RSTU} G_T^R(\tau) G_R^U(-\tau) V_R^S G_U^S(\tau) G_S^T(-\tau) V_T^U, \quad (33)$$

where the imaginary time dependent propagator matrices are given in terms of the decomposed factor orbitals Π_r^R and the Coulomb factors Λ_F^R :

$$G_S^R(\tau \geq 0) := + \sum_a \Pi_S^{*a} \Pi_a^R e^{-(\varepsilon_a - \mu)\tau} \quad (34)$$

$$G_S^R(\tau < 0) := - \sum_i \Pi_S^{*i} \Pi_i^R e^{-(\varepsilon_i - \mu)\tau} \quad (35)$$

$$V_R^S := \Lambda_R^{*F} \Lambda_F^S, \quad (36)$$

where μ is the Fermi level energy at zero temperature. The imaginary time integration can be done numerically on a small grid, for instance the minimax grid²⁹, which allows an evaluation of the MP2 correlation energy in $\mathcal{O}(N^4)$ in principle. In practice, however, this approach outperforms canonical MP2 calculations only for very large systems. The propagator matrices $G_S^R(\pm\tau)$ are analogous to the one body particle/hole Green's functions in real space and imaginary time $G_0(\mathbf{x}, \mathbf{x}'; \pm\tau)$. The time independent matrix propagator V_R^S corresponds to the Coulomb kernel. The factorization of the one

body energies $\varepsilon_{a,i}$ in imaginary time is equivalent to the Laplace transformed MP2 ansatz of Almlöf³⁰. Note that the imaginary time dependent matrices defined above are actually not propagators in the sense that

$$G_T^S(\tau_1) G_S^R(\tau_2) = G_T^R(\tau_1 + \tau_2) \quad (37)$$

for all R, T and $\tau_1, \tau_2 > 0$ for particles as well as $\tau_1, \tau_2 \leq 0$ for holes. They can only be used to directly connect vertices of the Coulomb interaction. If propagators in the above sense are required, one needs to employ a stricter ansatz for the factorization of the Coulomb vertex, namely

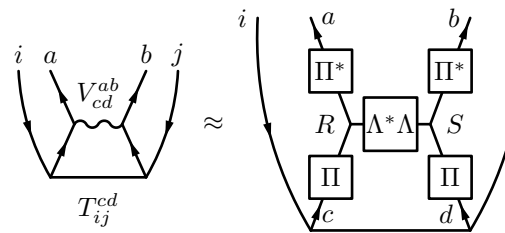
$$\Gamma_{rF}^q \approx \Lambda_F^R \Pi^{+q} \Pi_r^R, \quad (38)$$

where Π^+ denotes the Moore–Penrose pseudo inverse of Π . The convergence behavior of this ansatz remains, however, to be studied.

One can also evaluate the perturbation terms stochastically, directly using the real space Green's functions $G_0(\mathbf{x}, \mathbf{x}'; \pm\tau)$ rather than using the low rank propagators $G_S^R(\pm\tau)$. This has been done for one dimensional solids^{31,32} and for three dimensional solids³³.

B. Reduced scaling coupled cluster theory

The most demanding step in the canonical CCD (DCD) method using a plane wave basis set is the calculation of the particle/particle ladder contribution $T_{ij}^{cd} V_{cd}^{ab}$ in the amplitude equation, scaling as $\mathcal{O}(N_o^2 N_v^4)$ in time. Here, the factorized form of the Coulomb integrals V_{cd}^{ab} can be exploited to break down the simultaneous contraction over the indices c and d into a sequence of contractions involving only at most one index, as can be seen from the wiring diagram of the involved tensors:



$$\approx \quad (39)$$

The most expensive term in this sequence of contractions leads to a scaling of $\mathcal{O}(N_o^2 N_v N_R^2)$ in time, without exceeding the memory complexity of the coupled cluster amplitudes. As will be demonstrated in Section V, we find N_R to be proportional to the system size N , resulting in an $\mathcal{O}(N^5)$ scaling behavior in time of the particle/particle ladder contribution. Furthermore, the DCD amplitudes equations can be solely reformulated in an $\mathcal{O}(N^5)$ implementation with the use of the Coulomb vertex and its decomposed approximation, due to the absence of exchange terms between different clusters. Likewise, the most expensive term in CCSD (DCSD) amplitude equations includes the singles contribution to the

particle/particle ladder diagram ($T_{ij}^{cd}V_{ck}^{ab}T_d^k$, $T_{ij}^{cd}V_{ld}^{ab}T_c^l$, $T_{ij}^{cd}V_{kl}^{ab}T_d^l$). Similarly, these terms can be evaluated via the factor orbitals and the Coulomb factors in an $\mathcal{O}(N^5)$ scaling in time and in the DCSD approximation no term exceeds this scaling behavior.

V. RESULTS EMPLOYING THE LOW RANK FACTORIZATION

A. Total energies of the LiH solid

We first seek to discuss the convergence of the low rank factorization with respect to the number of iterations, N_R and the system size. To this end we study different supercell sizes constructed from two atomic LiH crystal unit cells including $2 \times 2 \times 2$, $3 \times 2 \times 2$, $3 \times 3 \times 2$, $4 \times 3 \times 2$ and $3 \times 3 \times 3$, corresponding to 16, 24, 36, 48 and 54 atoms, respectively. In this subsection we only employ MP2 theory to investigate the behavior of the correlation energy calculated from the factorized Coulomb integrals. This study focuses on the decomposition of the Coulomb vertex, neither employing the optimized auxiliary field nor the pseudized Gaussian type virtual orbitals³⁴ technique. The resulting Coulomb vertices that need to be fit are large, such that only the particle/hole part Γ_{iF}^a is used for this study. The kinetic energy cutoff defining N_F was set to 200 eV. The Li $2s^1$ and H $1s^1$ states have been treated as valence states.

Fig. 1 shows the relative error of the MP2 correlation energy retrieved as a function of the number of iterations. The relative error is computed from MP2 energies that employ integrals that have been calculated with and without the low rank tensor factorization. Fig. 1 reveals that the rate of convergence for the relative error of the MP2 energy is very similar in all different system sizes. We note that the smallest supercell containing 16 atoms only exhibits a slightly faster convergence. From these results we conclude that the required number of iterations in the tensor factorization algorithm is system size independent for intensive properties. Furthermore we find that 100 iterations are sufficient to achieve a relative accuracy of 1% with a rank that corresponds to $2N_F$.

The right side of Fig. 1 explores the convergence of the relative error in the MP2 correlation energy error retrieved as a function of N_R/N_F . Note that N_F corresponds to the number of plane wave vectors and scales linearly with respect to system size. This plot shows that the required rank N_R needed to achieve a certain relative level of accuracy also scales linearly with respect to the system size. Furthermore we find that systematically improvable exponential convergence can be achieved for this system and property by increasing N_R/N_F .

The computational cost for obtaining the TRD of the Coulomb vertex with $N_o = 27$, $N_v = 8469$, $N_F = N_G = 1830$ and $N_R = 1830$ is roughly 1000 CPU hours. Therefore the computational cost of the TRD exceeds the computational cost of a full MP2 calculation which is roughly

10 CPU hours in the present case despite the fact that the TRD formally scales more favorably with system size. However, we note that the present TRD algorithm is sufficiently efficient to reduce the total computational cost of coupled cluster theory calculations as discussed in the following.

B. Molecular adsorption of water on hexagonal boron nitride

We now turn to the application of the newly developed methodologies to some more challenging problems. We calculate the interaction between a water molecule and a hexagonal boron nitride (*h*BN) monolayer. We employ periodic coupled cluster doubles (CCD) and examine to what extent the TRD and the optimized auxiliary field approximations are accurate and efficient. We used the structures obtained by Al-Hamdani *et. al.*³⁵, whereby the molecule is oriented on top of an N site and the geometry has been optimized using the optB86b-vdW functional. The water–N distance was set to 3.2 Å. The *h*BN monolayer is modeled by 32 atoms in the periodic cell and the distance between two BN sheets was set to 16 Å. After checking convergence, we employed a 500 eV kinetic energy cutoff for the one particle orbitals along with Γ point sampling of the Brillouin zone. The B $2s^22p^1$, N $2s^22p^3$, O $2s^22p^4$, and H $1s^1$ states have been treated as valence states. Occupied HF states were converged within the full plane wave basis, whereas the virtual orbitals were constructed using Dunning’s contracted aug-cc-pVDZ (AVDZ) and aug-cc-pVTZ (AVTZ)^{36,37} pseudized Gaussians in a plane wave representation, orthogonalized to the occupied space³⁴. Pseudized Gaussians have proven to work reliably and efficiently for surface studies on periodic systems³⁴. Counterpoise corrections to the basis set superposition error (BSSE) were included in all correlated calculations. The adsorption energy is defined as the difference in energy between the noninteracting fragments and the interacting system

$$E_{\text{ads}} = E_{\text{H}_2\text{O}} + E_{\text{BN}} - E_{\text{H}_2\text{O}+\text{BN}}. \quad (40)$$

We note that in Ref. 35 the adsorption energy has been calculated as the difference between the total energy of water and *h*BN at the largest possible oxygen–surface distance of 8 Å and the total energy of water and *h*BN at the adsorption oxygen–surface distance.

Initially, we investigate the convergence of the adsorption energy with respect to the number of momentum grid points N_G , employed to evaluate the Coulomb vertex Γ_{rG}^q according to Eq. (7). The selection of the plane waves vectors \mathbf{G} is determined by a kinetic energy cutoff E_χ such that

$$\frac{\hbar^2 \mathbf{G}^2}{2m_e} < E_\chi. \quad (41)$$

For this purpose we utilize the pseudized AVDZ basis set for the virtual orbitals. The current system consists of




Figure 1. On the left the convergence of the relative error of the MP2 energy for $N_R = 2N_F$ with respect to fit iterations for different system sizes is shown. The convergence with respect to the rank N_R using 120 iterations is given on the right.

$N_o = 68$ and $N_v = 780$ occupied and virtual orbitals respectively. Kinetic energy cutoff values from 100 to 300 eV were employed for the calculation of the adsorption energy. The results are shown in the inset of Fig. 2. The adsorption energy behaves as $CE_\chi^{-3/2}$ ^{38,39} up to 200 eV, however, at higher cutoffs one observes a plateau in the curve as a result of the truncation of the virtual orbital space via the pseudized Gaussian basis functions. We conclude that a cutoff energy of 200 eV is sufficient to converge the adsorption energy to within 2–3 meV. We then investigate how accurately can the optimized auxiliary field approximate the Coulomb vertex $\tilde{\Gamma}_{rG}^q$ for the current system at the level of CCD theory. First we obtain the optimized Coulomb vertex $\Gamma_{rF}^q = U_G^{*F} \tilde{\Gamma}_{rG}^q$, where U_G^F consists of the left singular values of $\tilde{\Gamma}_{rG}^q$ associated to the N_F largest singular values, according to Eq. (8). The optimized Coulomb vertex is expected to be efficient since most of the space in the simulation cell is vacant, and the plane wave auxiliary basis contains redundant information. Different number of field variables N_F were employed to approximate the plane wave vectors N_G for the various cutoff energies. The behavior of the adsorption energy with respect to the number of field variables is shown in Fig. 2. The rapid convergence of the energy with increasing number of field variables owes to the locality of the molecular orbitals in the supercell. The adsorption energy obtained with a cutoff of 200 eV can be calculated within 0.5 meV accuracy using $N_F = 1450$ field variables to approximate $N_G = 4504$ plane wave vectors.

Since we can conclude that the adsorption energy can be computed within approximately 3 meV using a cutoff energy of 200 eV for the auxiliary plane wave basis and $N_F = 1450$ field variables to construct the optimized Coulomb vertex, we chose these settings to assess the accuracy of the low rank factorization of the Coulomb integrals. We used different number of vertex indices R to compute the factor orbitals and Coulomb factors of the Coulomb vertex, following Eq. (13). We approximate the particle/particle ladder contribution $T_{ij}^{cd} V_{cd}^{ab}$ in the am-

plitude equation of CCD via the factor orbitals and the Coulomb factors as shown in Eq. (39). The adsorption energy versus the number of the decomposed vertex indices N_R is shown in Fig. 3. The energy does not converge monotonically as in the case of the optimized auxiliary field approximation. This is due to the nonlinear nature of the canonical polyadic decomposition of the Coulomb vertex and the random initial choice for its factors. Nevertheless we observe a converged behavior with increasing N_R . Furthermore, a value of $N_R = 4350 = 3N_F$ is sufficient to yield an adsorption energy within 1 meV accuracy. This suggests that the TRD of the Coulomb vertex is a controllable approximation that can yield increasingly accurate results with increasing decomposition rank N_R . In order to further validate the accuracy of the TRD method we show the convergence of the absolute energy of the interacting system with respect to the decomposition rank N_R in the inset of Fig. 3. We observe an exponential convergence of the total energy. An accuracy better than 0.1% is achieved already with $N_R = 2N_F$. Nevertheless we stress that the corresponding accuracy in the adsorption energy is a result of an error cancellation of one to two orders of magnitude.

Having assessed the accuracy of the TRD we now calculate the adsorption energy of the water molecule on *h*BN using the AVTZ pseudized Gaussian basis set. The evaluation involves the decomposition of a Coulomb vertex with $N_o = 68$, $N_v = 1564$, $N_F = 0.33N_G = 1450$, and $N_R = 3N_F = 4350$. The computational cost to obtain the decomposed matrices is roughly 3000 CPU hours with 256 iterations. The results of the adsorption energy are shown in Table I. In order to grasp a physical insight of the system, we compare the CCD results with RPA+SOSEX and MP2 calculations⁴⁰. MP2 theory usually overestimates dispersion driven interactions, although in the description of BN bilayer interaction is fortuitously accurate⁴¹. Consequently, one expects MP2 theory to slightly overestimate the adsorption energy, whereas RPA+SOSEX is likely to yield a very accurate estimate. It is not surprising that CCD




Figure 2. Optimized auxiliary field (OAF) approximation error $\epsilon^{\text{OAF}}(N_F) = E_{\text{ads}}(N_F) - E_{\text{ads}}(N_G)$ of the CCD adsorption energy as a function of the number of field variables N_F used to approximate the Coulomb vertex. The number of G vectors N_G of the unapproximated Coulomb vertex $\tilde{\Gamma}_{rG}^q$ is indicated by the shape of the markers. The inset shows which marker corresponds to which N_G and plots the convergence of the adsorption energy with respect to N_G , corresponding to kinetic energy cutoff values of 100, 150, 200, 250, and 300 eV, respectively.

Table I. Adsorption energies of water on $h\text{BN}$ obtained using the pseudized Gaussian basis sets at different levels of theory. RPA+SOSEX calculations were performed using DFT PBE orbitals as reference, whereas MP2 and CCD using HF ones.

Basis set	RPA+SOSEX	MP2	CCD
AVDZ	62	83	54
AVTZ	72	92	62
AV(D,T)Z	76	95	65

underbinds the water molecule, since there exist findings that indicate the inability of CCD for an accurate description. Higher levels of theories, such as inclusion of the single excitations and the perturbative triples, are required for a more appropriate treatment. Nevertheless, the purpose of the current work is to examine the accuracy and efficiency of the newly developed methodologies rather than the accuracy of the method itself. The CPU hours required for the CCD calculations obtained with and without the TRD technique are summarized in Table II. The time for the evaluation of the particle/particle ladder term per iteration is as much as 43 times faster using a decomposition with $N_R = 2N_F$ and 22 times faster with $N_R = 3N_F$. This constitutes a significant gain in the computational effort of coupled cluster methods with only slight compromise in accuracy.




Figure 3. Low rank approximation error $\epsilon_{\text{ads}}^{\text{TRD}}(N_R) = E_{\text{ads}}(N_F, N_R) - E_{\text{ads}}(N_F)$ of the CCD adsorption energy as a function of the rank N_R using the OAF approximated Coulomb vertex with $N_F = 1450$ field variables. The inset shows the respective approximation error for surface+olecule fragment $\epsilon_{\text{H}_2\text{O}+\text{BN}}^{\text{TRD}} = E_{\text{H}_2\text{O}+\text{BN}}(N_F, N_R) - E_{\text{H}_2\text{O}+\text{BN}}(N_F)$ only, revealing an error cancellation of one to two orders of magnitude in the adsorption energy.

Table II. CPU hours per iteration comparing CCD calculation with and without the factorized Coulomb integrals. In parenthesis we denote the part for evaluating the particle/particle ladder term.

Basis set	$N_R = 2N_F$	$N_R = 3N_F$	no TRD
AVDZ	39 (13)	49 (24)	100 (75)
AVTZ	259 (28)	258 (55)	^a 1443 (1212)

^a estimation based on the AVDZ basis set.

VI. CONCLUSIONS

In this work we have outlined an algorithm to obtain a low rank tensor approximation of the Coulomb integrals having the same algebraic structure as its definition from the molecular orbitals:

$$V_{sr}^{pq} = \iint_{\mathbf{d}\mathbf{x} \mathbf{d}\mathbf{x}'} \psi^{*p}(\mathbf{x}) \psi^{*q}(\mathbf{x}') \frac{1}{|\mathbf{r} - \mathbf{r}'|} \psi_r(\mathbf{x}') \psi_s(\mathbf{x})$$

$$\approx \sum_{RS} \Pi_R^{*p} \Pi_S^{*q} \Lambda_R^{*F} \Lambda_F^S \Pi_r^S \Pi_s^R. \quad (42)$$

The factorization is obtained by fitting $\Lambda_F^S \Pi_S^{*q} \Pi_r^S$ to auxiliary three index quantities referred to as Coulomb vertices that are calculated from a resolution of identity approach using a plane wave basis set. In this manner the scaling of the computational cost for obtaining the low rank tensor approximation with respect to system size does not exceed $\mathcal{O}(N^4)$. To reduce the prefactor of the computational cost further we have outlined an approach to further compactify the representation of the Coulomb

vertices. We linearly transform the momentum index of the Coulomb vertices into a (truncated) basis referred to as an optimized auxiliary field. The accuracy of this truncation is systematically improvable using a single parameter that is used for the truncation of a singular value decomposition. The tensor factorization of the transformed Coulomb vertices is achieved using a regularized alternating least squares algorithm that converges rapidly using about 10^2 iterations only. In contrast, the nonregularized alternating least squares algorithm would require 10^5 – 10^6 iterations. We stress that we employ no prior assumptions for the real space grids used for expanding the low order tensors.

Once obtained, the tensor factorization of the Coulomb integrals can be employed to reduce the scaling of the computational cost of distinguishable coupled cluster theory to $\mathcal{O}(N^5)$ without further approximations. We demonstrate that the factorization can also be used to reduce the computational cost for evaluating the computationally most expensive term (particle/particle ladder diagram) in the coupled cluster doubles amplitude equations for the case of water adsorption on the *h*BN monolayer system. For system sizes containing 136 electrons in 1632 orbitals we achieve substantial reductions in the computational cost that are on the order of a factor 10–20 without compromising the accuracy and introducing any further approximation.

Future work will focus on combining the outlined techniques with explicitly correlated methods and finite size corrections in order to significantly expand the scope of periodic coupled cluster theory calculations using plane wave basis sets for solid state systems^{42,43}.

ACKNOWLEDGMENTS

We want to acknowledge fruitful discussions with Edgar Solomonik and Alexander Auer.

REFERENCES

- ¹A. Szabo and N. S. Ostlund, *Modern quantum chemistry: introduction to advanced electronic structure theory*. Mineola, N.Y.: Dover Publications, 1996.
- ²S. Reine, T. Helgaker, and R. Lindh, “Multi-electron integrals,” 2012.
- ³A. Hellweg, C. Hättig, S. Höfener, and W. Klopper, “Optimized accurate auxiliary basis sets for RI-MP2 and RI-CC2 calculations for the atoms Rb to Rn,” *Theor. Chem. Acc.*, vol. 117, no. 4, pp. 587–597, 2007.
- ⁴R. Ahlrichs, “Efficient evaluation of three-center two-electron integrals over Gaussian functions,” *Phys. Chem. Chem. Phys.*, vol. 6, no. 22, pp. 5119–5121, 2004.
- ⁵F. Weigend and R. Ahlrichs, “Balanced basis sets of split valence{,} triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy,” *Phys. Chem. Chem. Phys.*, vol. 7, no. 18, pp. 3297–3305, 2005.
- ⁶F. R. Manby and P. J. Knowles, “Poisson Equation in the Kohn-Sham Coulomb Problem,” *Phys. Rev. Lett.*, vol. 87, p. 163001, sep 2001.
- ⁷A. Grüneis, M. Marsman, J. Harl, L. Schimka, and G. Kresse, “Making the random phase approximation to electronic correlation accurate,” *J. Chem. Phys.*, vol. 131, no. 15, p. 154115, 2009.
- ⁸D. Kats, “The distinguishable cluster approach from a screened Coulomb formalism,” *The Journal of Chemical Physics*, vol. 144, p. 044102, Jan. 2016.
- ⁹D. Kats and F. R. Manby, “Communication: The distinguishable cluster approximation,” *The Journal of Chemical Physics*, vol. 139, p. 021102, July 2013.
- ¹⁰D. Kats, D. Kreplin, H.-J. Werner, and F. R. Manby, “Accurate thermochemistry from explicitly correlated distinguishable cluster approximation,” *The Journal of Chemical Physics*, vol. 142, p. 064111, Feb. 2015.
- ¹¹E. G. Hohenstein, R. M. Parrish, and T. J. Martínez, “Tensor hypercontraction density fitting. I. Quartic scaling second- and third-order Møller-Plesset perturbation theory,” *The Journal of Chemical Physics*, vol. 137, p. 044103, July 2012.
- ¹²R. M. Parrish, E. G. Hohenstein, T. J. Martínez, and C. D. Sherrill, “Tensor hypercontraction. II. Least-squares renormalization,” *The Journal of Chemical Physics*, vol. 137, p. 224106, Dec. 2012.
- ¹³E. G. Hohenstein, R. M. Parrish, C. D. Sherrill, and T. J. Martínez, “Communication: Tensor hypercontraction. III. Least-squares tensor hypercontraction for the determination of correlated wavefunctions,” *The Journal of Chemical Physics*, vol. 137, p. 221101, Dec. 2012.
- ¹⁴N. Shenvi, H. v. Aggelen, Y. Yang, W. Yang, C. Schwerdtfeger, and D. Mazziotti, “The tensor hypercontracted parametric reduced density matrix algorithm: Coupled-cluster accuracy with $\mathcal{O}(r4)$ scaling,” *The Journal of Chemical Physics*, vol. 139, p. 054110, Aug. 2013.
- ¹⁵U. Benedikt, K.-H. Böhm, and A. A. Auer, “Tensor decomposition in post-Hartree-Fock methods. II. CCD implementation,” *The Journal of Chemical Physics*, vol. 139, p. 224101, Dec. 2013.
- ¹⁶J. Harl, *The linear response function in density functional theory*. PhD thesis, University of Vienna, Vienna, 2008.
- ¹⁷P. E. Blöchl, “Projector augmented-wave method,” *Physical Review B*, vol. 50, pp. 17953–17979, Dec. 1994.
- ¹⁸G. Kresse and J. Hafner, “Norm-conserving and ultrasoft pseudopotentials for first-row and transition elements,” *Journal of Physics: Condensed Matter*, vol. 6, no. 40, p. 8245, 1994.
- ¹⁹G. Kresse and J. Furthmüller, “Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set,” *Physical Review B*, vol. 54, pp. 11169–11186, Oct. 1996.
- ²⁰R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, “Monte Carlo calculations of coupled boson-fermion systems. I,” *Phys. Rev. D*, vol. 24, pp. 2278–2286, Oct. 1981.
- ²¹F. L. Hitchcock, “The expression of a tensor or a polyadic as a sum of products,” *J. Math. Phys.*, vol. 6, no. 1, pp. 164–189, 1927.
- ²²T. G. Kolda and B. W. Bader, “Tensor Decompositions and Applications,” *SIAM Review*, vol. 51, pp. 455–500, Aug. 2009.
- ²³L. Karlsson, D. Kressner, and A. Uschmajew, “Parallel algorithms for tensor completion in the CP format,” *Parallel Computing*, 2015.
- ²⁴E. H. Moore, “The fourteenth western meeting of the American Mathematical Society,” *Bull. Amer. Math. Soc.*, vol. 26, no. 9, pp. 385–396, 1920.
- ²⁵R. Penrose and J. A. Todd, “A generalized inverse for matrices,” *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 51, p. 406, July 1955.
- ²⁶S. Kindermann and C. Navasca, “Analysis and approximation of the Canonical Polyadic tensor decomposition,” *arXiv preprint arXiv:1109.3832*, 2011.
- ²⁷A. Cichocki and R. Zdunek, “Regularized alternating least squares algorithms for non-negative matrix/tensor factorization,” in *Advances in Neural Networks—ISNN 2007*, pp. 793–802, Springer, 2007.
- ²⁸N. Li, S. Kindermann, and C. Navasca, “Some Convergence Results on the Regularized Alternating Least-Squares Method for Tensor Decomposition,” *arXiv:1109.3831 [math]*, Sept. 2011. arXiv: 1109.3831.

- ²⁹M. Kaltak, J. Klimeš, and G. Kresse, “Low Scaling Algorithms for the Random Phase Approximation: Imaginary Time and Laplace Transformations,” *J. Chem. Theory Comput.*, vol. 10, pp. 2498–2507, June 2014.
- ³⁰J. Almlöf, “Elimination of energy denominators in Møller–Plesset perturbation theory by a Laplace transform approach,” *Chemical Physics Letters*, vol. 181, pp. 319–320, June 1991.
- ³¹S. Y. Willow, K. S. Kim, and S. Hirata, “Stochastic evaluation of second-order many-body perturbation energies,” *The Journal of Chemical Physics*, vol. 137, p. 204122, Nov. 2012.
- ³²D. Neuhauser, R. Baer, and D. Zgid, “Stochastic self-consistent Green’s function second-order perturbation theory (sGF2),” *arXiv:1603.04141 [cond-mat, physics:physics]*, Mar. 2016. arXiv: 1603.04141.
- ³³T. Schäfer and G. Kresse. private communication.
- ³⁴G. H. Booth, T. Tsatsoulis, G. K.-L. Chan, and A. Grüneis, “From plane waves to local Gaussians for the simulation of correlated periodic systems,” *The Journal of Chemical Physics*, vol. 145, p. 084111, Aug. 2016.
- ³⁵Y. S. Al-Hamdani, M. Ma, D. Alfè, O. A. v. Lilienfeld, and A. Michaelides, “Communication: Water on hexagonal boron nitride from diffusion Monte Carlo,” *The Journal of Chemical Physics*, vol. 142, p. 181101, May 2015.
- ³⁶T. H. Dunning Jr., “Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen,” *The Journal of Chemical Physics*, vol. 90, pp. 1007–1023, Jan. 1989.
- ³⁷D. Feller, “The role of databases in support of computational chemistry calculations,” *Journal of Computational Chemistry*, vol. 17, pp. 1571–1586, Oct. 1996.
- ³⁸J. Harl and G. Kresse, “Cohesive energy curves for noble gas solids calculated by adiabatic connection fluctuation-dissipation theory,” *Physical Review B*, vol. 77, p. 045136, Jan. 2008.
- ³⁹M. Marsman, A. Grüneis, J. Paier, and G. Kresse, “Second-order Møller–Plesset perturbation theory applied to extended systems. I. Within the projector-augmented-wave formalism using a plane wave basis set,” *The Journal of Chemical Physics*, vol. 130, p. 184103, May 2009.
- ⁴⁰J. Paier, X. Ren, P. Rinke, G. E. Scuseria, A. Grüneis, G. Kresse, and M. Scheffler, “Assessment of correlation energies based on the random-phase approximation,” *New Journal of Physics*, vol. 14, no. 4, p. 043002, 2012.
- ⁴¹F. Hummel, T. Gruber, and A. Grüneis, “A many-electron perturbation theory study of the hexagonal boron nitride bilayer system*,” *The European Physical Journal B*, vol. 89, no. 11, p. 235, 2016.
- ⁴²A. Grüneis, “Efficient Explicitly Correlated Many-Electron Perturbation Theory for Solids: Application to the Schottky Defect in MgO,” *Phys. Rev. Lett.*, vol. 115, p. 66402, aug 2015.
- ⁴³K. Liao and A. Grüneis, “Communication : Finite size correction in periodic coupled cluster theory calculations of solids Communication : Finite size correction in periodic coupled cluster theory calculations of solids,” *J. Chem. Phys.*, vol. 141102, no. 15, pp. 1–5, 2016.